

METHOD FOR DISCOVERY AND ROUTING WITHIN MOBILE AD-HOC NETWORKS

RELATED APPLICATION DATA

This application is related to U.S. Patent Application Serial No. 09/773,682, titled “ENABLING RESTRICTED COMMUNICATIONS BETWEEN A PLURALITY OF USERS”, filed January 31, 2001, and to U.S. Patent Application Serial No. 09/930,779, titled “METHOD FOR DISCOVERY AND ROUTING USING BUDDY LISTS IN MOBILE AD-HOC NETWORKS”, filed August 15, 2001, and to U.S. Patent Application Serial No. 09/948,300, titled “METHOD FOR RESTRICTED COMMUNICATIONS USING BUDDY LISTS IN MOBILE AD-HOC NETWORKS”, filed September 6, 2001, and to U.S. Patent Application Serial No. 09/948,270, titled “METHOD FOR DISCOVERY AND ROUTING USING ATTRIBUTES IN MOBILE AD-HOC NETWORKS”, filed September 6, 2001, and to U.S. Patent Application Serial No. _____, titled “METHOD FOR DISCOVERY AND ROUTING USING A PRIORI KNOWLEDGE IN THE FORM OF APPLICATION PROGRAMME WITHIN MOBILE AD-HOC NETWORKS”, filed _____.

FIELD OF THE INVENTION

This invention pertains to mobile ad-hoc networks, and more particularly to communication within a mobile ad-hoc network.

BACKGROUND OF THE INVENTION

The ability to communicate with others has always been an important function for people. First the telegraph, and then the telephone revolutionized communication, by making communication at a distance possible. The cellular telephone marked a further turning point in communication, as people were no longer tied to a phone with a fixed physical location. And the more recent Bluetooth standard (Specification of the Bluetooth System, Version 1.1, February 22, 2001) and 802.11a standard (IEEE Standard 802.11a-1999, available on the Institute of Electrical and Electronics Engineers web site, <http://standards.ieee.org/reading/ieee/std/lanman/802.11a-1999.pdf>) describe how wireless communication may be achieved using mobile devices.

Where mobile devices attempt to intercommunicate, a mobile ad-hoc network (MANET) is set up. In contrast to the typically-envisioned computer network, a MANET does not consist of separate servers and clients. Instead, each mobile device acts as both client and server. That is, each mobile device is responsible for not only sending messages on its own, but also for routing messages. Further, each mobile device needs to be able to route messages sent by a different device, and bound for a different destination.

FIG. 1A shows a scattering of mobile ad-hoc network devices 105. In FIG. 1A, a number of mobile devices are scattered around a location. The location may be a mall, sporting arena, city, or any other type of location, without any limit. In the discussion below, for simplicity of description each person using a mobile device will be described as able to communicate with other persons, although in truth it is the mobile devices that communicate with each other. For example, although the description below may talk about Aaron 110 communicating with Beth 125, it is really Aaron's device that communicates with Beth's device.

Each mobile device has a limited range: for example, Aaron's device 110 has a range 115 shown as a dashed circle. Thus, not every other mobile device has a direct line of communication with every other device. For example, Beth 125 is within direct communications range of Aaron 110, but Charlie 130 is not within direct communications range of Aaron 110.

The intuitive approach for managing a MANET is to have each device be aware of all other devices in the MANET. FIG. 1B shows a mobile ad-hoc network 195 established among the mobile ad-hoc network devices of FIG. 1A. As may be seen in FIG. 1B, Aaron 110 may communicate directly with Beth 125, Elizabeth 140, Harry 155, and Isis 160, but may not communicate directly with Charlie 130 or David 135. Nevertheless, Aaron 110 may communicate with Charlie 130 or David 135 through Beth 125 or Isis 160, respectively.

Even if a person is in the MANET, he may not be reachable. For example, consider Oscar 190. Even if Oscar's device is active within the network, given Oscar's location, Oscar 190 may communicate with other devices in the MANET only through Mark 180. If Mark 180 leaves the network (e.g., by turning off his device), Oscar 190 may neither send nor receive messages from anyone else in the MANET.

The intuitive approach to managing a MANET imposes global limitations on the MANET, which may affect conditions at a local level. For example, devices in the MANET

may communicate only with one device at a time. Therefore, devices have to schedule their communications with neighboring devices.

The present invention addresses this and other problems associated with the prior art.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1A shows a typical arrangement of mobile ad-hoc network devices.

FIG. 1B shows a mobile ad-hoc network established among the mobile ad-hoc network devices of FIG. 1A.

FIG. 2 shows a mobile device operable within the mobile ad-hoc network of FIG. 1B, according to an embodiment of the invention.

FIG. 3A shows a timing diagram for synchronization of the mobile devices in the mobile ad-hoc network of FIG. 1B before the introduction of Aaron's device, according to an embodiment of the invention.

FIG. 3B shows a timing diagram for synchronization of the mobile devices in the mobile ad-hoc network of FIG. 1B after the introduction of Aaron's device, according to an embodiment of the invention.

FIG. 4 shows a timing diagram for data transfer between two mobile devices in the mobile ad-hoc network of FIG. 1B, according to an embodiment of the invention.

FIG. 5 shows the mobile device of FIG. 2 including a synchronization table, according to an embodiment of the invention.

FIG. 6 shows the procedure used by the mobile device of FIG. 2 to synchronize with other mobile devices in the mobile ad-hoc network of FIG. 1B, according to an embodiment of the invention.

FIGs. 7A and 7B show the procedure used by the mobile device of FIG. 2 to transfer data to another mobile device in the mobile ad-hoc network of FIG. 1B, according to an embodiment of the invention.

DETAILED DESCRIPTION

Referring momentarily to FIG. 1B, the discussion below only refers to sub-network 198 consisting of users Aaron 110, Beth 125, David 135, Elizabeth 140, Frank 145, and Joshua 165. A person skilled in the art will recognize how the embodiment of the invention may be extended to the entire mobile ad-hoc network (MANET) 195 shown.

FIG. 2 shows a mobile device operable within the mobile ad-hoc network (MANET) of FIG. 1B, according to an embodiment of the invention. In FIG. 2, mobile device 205 is used for communicating with other nearby mobile devices. Although shown as a personal digital assistant (PDA), a person skilled in the art will recognize that mobile device 205 may be any device capable of interacting with nearby devices, typically in a wireless manner. For example, mobile device 205 may be a cellular telephone or notebook computer properly equipped to interact with nearby devices. In fact, mobile device 205 need not be truly mobile: device 205 might be a fixed access point (such as a desktop computer or kiosk within the environment). But for purposes of the discussion below, device 205 will be considered to be mobile.

Mobile device 205 includes, among other components, receiver 210, transmitter 215, synchronization system 220, and memory 225. Receiver 210 is used to receive communications from other nearby mobile devices. Transmitter 215 is used to transmit communications to other nearby mobile devices. Although in a preferred embodiment receiver 210 and transmitter 215 are designed to operate wirelessly, a person skilled in the art will recognize that receiver 210 and transmitter 215 may also be wire-based. Synchronization system 220 is to synchronize device 205 with other nearby devices. And memory 225 is used to store information of a temporary or semi-permanent nature within mobile device 205.

Synchronization system 220 includes three components. First is arbitrator 230. Arbitrator 230 is used to arbitrate the timing and manner of synchronizing with other nearby devices. Second is synchronization table updater 235. Synchronization table updater 235 updates the synchronization table (discussed below) with data as arbitrated by arbitrator 230. Finally, clock 240 is the internal clock of device 205, and is used to monitor the passage of time, as viewed by device 205.

In mobile device 205, memory 225 includes synchronization table 245. Synchronization table 245 stores information about the synchronization between device 205 and other nearby devices. Specifically, synchronization table 245 stores device identifiers, the frequencies with which device 205 communicates with other nearby devices, and the timing interval at which device 205 checks with other nearby devices to see if data is to be transferred.

Among the various operations managed by synchronization system 220 is the entry and exit of devices within the range of device 205. Specifically, as devices enter and leave,

the nearby devices with which mobile device 205 may communicate varies. Synchronization system 220 is responsible for determining changes in which devices are nearby and establishing a communications protocol (using arbitrator 230) with new nearby devices.

FIG. 3A shows a timing diagram for synchronization of the mobile devices in the mobile ad-hoc network of FIG. 1B before the introduction of Aaron's device, according to an embodiment of the invention. Time is represented horizontally. Similarly-hatched boxes in rows for two different devices represent communication between the two devices. For example, boxes 305 and 310 show that Frank and Joshua are synchronizing.

Note that different pairs of devices may synchronize at the same time. For example, as shown by boxes 315 and 317 and boxes 320 and 322, Beth and David may synchronize at the same time as Elizabeth and Frank.

FIG. 3B shows a timing diagram for synchronization of the mobile devices in the mobile ad-hoc network of FIG. 1B after the introduction of Aaron's device, according to an embodiment of the invention. Aaron has synchronized with the other devices, as shown by row 325. Aaron has agreed to frequencies and synchronization periods/intervals with Beth, Elizabeth, and Joshua, the devices within range of Aaron. For example, as shown by boxes 330 and 35, Aaron and Elizabeth have agreed to mutually-acceptable synchronization periods.

A few comments about synchronization periods are in order. First, the pairs of devices may arbitrate to any mutually-acceptable frequency, communications technology, and synchronization period. The frequency is the particular channel in the electromagnetic spectrum at which the devices communicate. The communications technology specifies how the frequency is to be used for communication. And the synchronization period identifies how frequently the devices synchronize with each other. No limitations are placed on the arbitration at all. In the event a proposed synchronization period collides (in frequency, communications technology, or timing of synchronization) with another for one of the devices, the pair may determine a different synchronization period.

As mentioned briefly above, the communications technology specifies how the devices use the frequency for communication. There are many different communications technologies that the devices may utilize. For example, the devices may agree to a frequency-hopping schedule, so that they change their frequency of communication every so often. Or the devices may use chipping code, spreading the communications across a range of frequencies using a key. Without the key to identify the proper signals, the

communication sounds like static. A person skilled in the art will recognize other communications technologies that may be used.

Note that it is possible for nearby pairs of devices to communicate using the same frequency, provided the pairs of devices are capable of differentiating the communications.

5 For example, two adjacent pairs of devices may both use chipping code at the same frequency, if their keys allow them to adequately separate the communications.

Second, the interval between synchronization periods may vary. Depending on the type of communication between the devices, either a low or high latency may be desired. For example, if the mobile devices are transmitting speech, low latency is desired, to avoid long
10 breaks in the spoken output. On the other hand, if one of the mobile devices is a sensor and the other mobile device is a system using the sensor to check a status, only an occasional update is needed, and high latency is acceptable.

Third, the length of the synchronization may vary. Proper synchronization may be more or less critical. By providing a longer synchronization period, a higher degree of certainty of synchronization may be achieved.

Fourth, different communication frequencies may be selected by the pairs of devices. If all devices had to use the same frequency, any signal impinging from one pair of devices might cause destructive interference to nearby pairs of communicating devices. But if pairs of devices can agree to different frequencies and/or communication technologies, the
20 destructive interference from one pair of devices will not impair other nearby devices.

Fifth, note that synchronization periods may be dynamic. Specifically, based on changes in circumstances, a pair of devices may begin a new arbitration and determine a new synchronization period. For example, consider the sensor referred to above. If the value detected by the sensor goes out of range, the sensor may need to report its value more
25 frequently. The sensor may rearbitrate its synchronization period with a neighboring device, to allow for more frequent reports. A person skilled in the art will recognize that other components of the synchronization period, such as frequency, communications technology, or interval, may also be changed as needed.

Sixth, a person skilled in the art will recognize that the duration of the
30 synchronization period is typically as short as possible. Devices engage in synchronization as quickly as possible, so as to spend as much time as possible actually transmitting network traffic, or to conserve power. Nevertheless, a person skilled in the art will recognize that the duration may also be negotiated, and can specifically be longer than the minimum needed.

Other factors may affect the arbitration of frequency, synchronization period, and time interval. For example, traffic patterns, data urgency, quality of service requirements, battery condition, and device life expectancy may affect the arbitration. A person skilled in the art may recognize other factors as well.

5 Note that during interactions for arbitration and synchronization, no message or application data is transmitted. Only arbitration and synchronization data is transmitted at these times. Synchronization is used to make sure that the devices' clocks are correlated. (Note that the clocks do not have to use the same time periods, only that the devices are aware of how to agree to a common time interval.) Data is transmitted at other times,
10 specifically between scheduled events like synchronization.

But before a device may transmit data, the device has to notify the intended recipient of the data about the communication. This notification is done during the synchronization period. Thus, if a device has data to transmit, the device waits until the synchronization period with the intended recipient device. The device then informs the intended recipient that there is data to transfer. At the first time interval open to both devices, data is transmitted.

FIG. 4 shows a timing diagram for data transfer between two mobile devices in the mobile ad-hoc network of FIG. 1B, according to an embodiment of the invention. In FIG. 4, Aaron has data to transmit to Joshua. During the synchronization period (boxes 402 and 403), Aaron notifies Joshua of the data to transmit. Then, at the next time interval when both devices are free, (boxes 405 and 410), Aaron transmits data to Joshua. After notifying Joshua of the data to transfer, both devices wait for their first mutual free interval. In FIG. 4, both devices are free immediately after their mutual synchronization period. So, as shown by boxed 405 and 410, Aaron begins to transfer data to Joshua.

Note that shortly after beginning transmission, Aaron is busy synchronizing with
25 another mobile device (shown by box 415). During this time, Aaron may not transfer data to Joshua (note that Joshua is idle during this time interval). But once Aaron and Joshua are both free again, Aaron resumes transmitting data to Joshua (as shown by boxes 420 and 425).

FIG. 5 shows the mobile device of FIG. 2, according to an embodiment of the invention. In FIG. 5, the details of synchronization table 505 are shown, assuming mobile
30 device 205 belongs to Aaron. Each entry in synchronization table 505 represents communication with another device. Note that synchronization table includes four pieces of data: an identifier for the device with which communication occurs, a frequency upon which communication is to occur, a communications technology to use in communicating with the

device, and a time interval at which the communication is to occur. Thus, entry 510 specifies the information used to synchronize with Elizabeth, entry 515 specifies the information used to synchronize with Frank, and entry 520 specifies the information used to synchronize with Joshua.

5 Time interval may be specified in a number of ways: for example, during each synchronization period, the two devices may agree to the time at which they will next communicate, in which case the agreed time (in terms of the local clock) is stored in synchronization table 505. Or, synchronization may occur at regular intervals, and the time interval may specify the time at which synchronization with that device first occurred (and
10 modulo arithmetic may be used to determine the next time of synchronization).

Note that in all of the above discussion, only a pair of devices is involved. The arbitration and synchronization of each pair of devices creates a peer-to-peer relationship between pairs of devices. The network as a whole is an emergent property of the peer-to-peer relationships between pairs of devices. For example, referring back to FIG. 1B, David 135 may not communicate directly with Frank 145. But if David 135 has a message for Frank 145, David 135 uses the synchronization period established between David 135 and Elizabeth 140 to transmit the message to Elizabeth 140. Elizabeth 140, in turn, uses the synchronization period established between Elizabeth 140 and Frank 145 to transmit the message to Frank 145. In this manner, the network is established, using the peer-to-peer relationships.

FIG. 6 shows the procedure used by the mobile device of FIG. 2 to synchronize with other mobile devices in the mobile ad-hoc network of FIG. 1B, according to an embodiment of the invention. In FIG. 6, at block 605, the device discovers a second device in range. At block 610, the devices communicate to discuss a synchronization period. Finally, at block
25 615, the devices arbitrate a frequency and time interval to use for synchronization.

FIGs. 7A and 7B show the procedure used by the mobile device of FIG. 2 to transfer data to another mobile device in the mobile ad-hoc network of FIG. 1B, according to an embodiment of the invention. In FIG. 7A, at block 705, the device waits for the synchronization period with the intended recipient. At block 710, the devices synchronize.
30 At block 715, the devices reset their local clocks. At decision point 720, the devices check to see if there is data to transmit to the other device. If there is no data to transmit, the device returns to block 705 to wait for the next synchronization period. (Typically, the device will be waiting to synchronize with another device, but the device may resume transmitting data

to or receiving data from a third device, if the communication was interrupted to synchronize with the other device.)

If there is data to transmit, then at block 725 (FIG. 7B) the device informs the other device of the data. At block 730, the devices wait for a mutually-free time interval. Finally,
5 at block 735 the data is transmitted, and the device returns to block 705 to wait for another synchronization event.

If there is no data to transmit or receive, the device may power down either partially (transceiver) or wholly (complete device) to reduce power consumption. The device would naturally power back up at the agreed-upon synchronization periods.

10 Having illustrated and described the principles of my invention in an embodiment thereof, it should be readily apparent to those skilled in the art that the invention can be modified in arrangement and detail without departing from such principles. I claim all modifications coming within the spirit and scope of the accompanying claims.